

PATCHWORK: A MULTIPATCH INFRASTRUCTURE FOR MULTIPHYSICS/MULTISCALE/MULTIFRAME FLUID SIMULATIONS

HOTAKA SHIOKAWA^{1,2}, ROSEANNE M. CHENG², SCOTT C. NOBLE³, JULIAN H. KROLIK²

¹Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA

² Department of Physics and Astronomy, Johns Hopkins University, Baltimore, MD 21218, USA
 and

³Department of Physics and Engineering Physics, University of Tulsa, Tulsa, OK 74104, USA

ABSTRACT

We present a “multipatch” infrastructure for numerical simulation of fluid problems in which sub-regions require different gridscales, different grid geometries, different physical equations, or different reference frames. Its key element is a sophisticated client-router-server framework for efficiently linking processors supporting different regions (“patches”) that must exchange boundary data. This infrastructure may be used with a wide variety of fluid dynamics codes; the only requirement is that their primary dependent variables be the same in all patches, e.g., fluid mass density, internal energy density, and velocity. Its structure can accommodate either Newtonian or relativistic dynamics. The overhead imposed by this system is both problem- and computer cluster architecture-dependent. Compared to a conventional simulation using the same number of cells and processors, the increase in runtime can be anywhere from negligible to a factor of a few; however, one of the infrastructure’s advantages is that it can lead to a very large reduction in the total number of zone-updates.

Keywords: methods:numerical — hydrodynamics — MHD

1. INTRODUCTION

1.1. Importance of multiphysics/multiscale/multiframe capability

Many important physical processes involve heterogeneous systems in which the nature of the matter in different regions exhibits strong contrasts. The material may vary in its characteristic internal length or time scales, or in its local geometric symmetry. There may even be contrasts in which the physical mechanisms of importance differ between regions: for example, chemical reactions or self-gravity may be significant in some, but not all locations. These regions may also move with respect to one another, perhaps changing shape as they do. When the regions have relative motion, the fact that physics is often most concisely described in a system’s mean rest-frame means that no single rest-frame is appropriate for the entire problem. At the same time, interactions between these regions may nonetheless demand simulation methods allowing data from one region to inform the behavior of another.

Problems exhibiting strong contrasts in length or time scales are called “multiscale problems”. We will also use this term to include contrasts in grid symmetry. In multiscale problems, numerical methods work best with

different grid systems in different regions, perhaps contrasting in resolution, perhaps in symmetry, e.g., polar vs. Cartesian. Those involving disparities in mechanisms are called “multiphysics problems”. In these problems, one must solve entirely different equations: those of magnetohydrodynamics (MHD) rather than those of hydrodynamics, or with or without transport processes such as viscosity or diffusion. Problems with internal frame shifts we dub “multiframe problems”. For these, it would be desirable to translate the equations from one frame to another in different portions of the calculation. Astrophysics is rich in problems to which at least one, and sometimes all, of these labels apply, and therefore at least one, and sometimes all, of the difficulties, both technical and conceptual, that they pose.

To illustrate their significance, consider a few examples. The topic that initially motivated our work is the mechanics of accretion around a binary system. For us, the partners in the binary are supermassive black holes (see, e.g., [Schnittman 2013](#)), but much about this problem changes little whether the binary comprises a pair of proto-stars (e.g., as imaged and analyzed by [Mayama et al. 2010](#)) or a pair of black holes. In this situation, there are widely disparate scales because the structure of

the circumbinary disk varies on the scale a of the orbital separation, whereas most of the accretion power emerges at the inner edges of the disks orbiting around the individual masses (often called “mini-disks”), which could be a great deal smaller. In addition, throughout these disks, even to define the saturation level of the MHD turbulence supplying the accretion torques requires treatment of lengthscales small compared to the disk scale heights, which could be considerably smaller than the radial scale. There are also differences in the symmetry of well-designed local grids. Because angular momentum transport in accretion disks is slow compared to the orbital time, it is very important that there be little numerical momentum diffusion; this fact demands a grid mimicking the symmetry of their nearly-circular flow. However, such a grid would be polar and centered on the binary center-of-mass for the circumbinary disk, whereas for each mini-disk, it would be polar and centered on the object whose gravity is most important for that disk. A single Cartesian grid for the entire system would likely produce an intolerable level of numerical diffusion. This binary accretion problem is also one that demands multiple reference frames for much the same reason it requires multiple sub-grids with different symmetries. The physics of the circumbinary flow is easiest to grasp in the center-of-mass frame; that of the individual mini-disks in the frame of each member of the binary. Thus, one would like to be able to divide this calculation into at least three different zones, each with its own grid and reference frame.

Another example may be found in tidal disruption of stars by supermassive black holes, which has become a subject of great interest in recent years as numerous examples have been found (in both optical/UV, e.g. [Gezari et al. \(2009\)](#) and [Arcavi et al. \(2014\)](#), and in X-rays: [Auchettl et al. \(2016\)](#)). This is a multiscale problem because it is necessary both to resolve dynamics within the star as it is broken apart and to follow the fluid dynamics of the debris as it gradually accretes onto the black hole. Measured in terms of gravitational radii r_g defined relative to the mass M of the black hole ($r_g \equiv GM/c^2$), main sequence stars are $\sim 1M_6^{-1}r_g$ in diameter, where M_6 is the black hole in units of $10^6 M_\odot$. Thus, to follow their break-up requires cells $\ll 0.1M_6^{-1}r_g$ in size. On the other hand, the debris orbits have semi-major axes $\sim 10^3 M_6^{-1/3}r_g$, so that the fluid motion after stellar break-up takes place on a much larger scale. Nonetheless, despite this dramatic scale contrast, the break-up of the star is inextricably tied to the much larger-scale debris motion. It is also a multiphysics problem because stellar self-gravity, not surprisingly, is of the essence so long as the star stays in one piece, but after its matter is spread sufficiently widely, it

becomes inconsequential. And it is a multiframe problem because the mechanics of a nearly hydrostatic star are definitely best viewed in the star’s frame where the fluid velocities are small, whereas the mechanics of an accretion flow are far more easily understood in the black hole frame. Its multiframe nature also creates a contrast in grid symmetry because coherent stars are most naturally treated in a spherical coordinate system whose origin is the center of the star, whereas flow around a black hole is best described in a cylindrical or spherical coordinate system whose origin is the center of the black hole. Thus, this problem, too, involves all these categories of complication.

1.2. State of the art and its limitations

The desirability of overcoming these challenges has not gone entirely unnoticed by the computational community, and a number of partial solutions have been developed. Adaptive Mesh Refinement (AMR) methods can dynamically adjust spatial resolution to follow local lengthscales ([Berger & Olinger 1984](#); [Berger & Colella 1989](#)); a closely-related scheme, overlapping moving grids, can be used to follow a coherent region with a distinct spatial scale or symmetry. Multiprogram/multiData (MPMD) methods ([Barney 2017](#)) offer a convenient way to evolve different regions according to the different mechanisms acting in them¹. We will call these separate regions “patches”, hence the name “multipatch” for our general approach. Because each patch is run by an independent program under MPMD, their only interaction is through the exchange of boundary conditions.

For some of these methods, professionally-supported implementations are available. *Chombo*, for example, is a particularly well-developed AMR package ([Adams et al. 2014](#)). It permits the use of two different methods to divide up regions into separate grids, embedded boundaries and mapped multiblocks. There are also relativistic versions of these fixed multiblock methods ([Clough et al. 2015](#); [Schnetter et al. 2014](#)). General relativistic hydrodynamics has been treated ([Blakely et al. 2015](#)) using the techniques of *Overture* ([Brown et al. 1997](#)), which offers users the option of moving overlapping grids. Numerical relativity calculations can use multiblock infrastructure with AMR, but all dependent variables must be defined to a global Cartesian tensor basis ([Pollney et al. 2017](#)). Similarly, there are numerous MPMD systems permitting computation of different physics in different parts of a global system. These

¹ In practical terms, users run multiple executables all sharing the same Message Passing Interface (MPI) “MPI.COMM.WORLD” communicator

range in their applications from linking multiscale fluid simulation to molecular dynamics (Nie et al. 2006) to modeling bloodflow through the brain (Grinberg et al. 2013).

Another approach to solving the problems of multiple scales, but not multiple physics, is the use of moving unstructured grids (e.g., the codes *AREPO*: Springel 2010 and *TESS*: Duffell & MacFadyen 2011). Schemes like these very flexibly place resolution where it is required for the hydrodynamics. It has also recently become possible to extend them from hydrodynamics to magnetohydrodynamics (Duffell 2016; Mocz et al. 2016). They do not, however, naturally retain the virtues of conforming to natural symmetries of the problem (e.g., suppressing numerical diffusion by aligning cell axes with the fluid velocity), nor do they readily permit the use of contrasting physics in different regions. With significant effort, it is possible to avoid the first drawback (Duffell 2016), but a new solution must be created for each new problem.

Despite the real successes of all these different schemes, there remain significant barriers to their employment on many kinds of problems. Multiblocks must fit smoothly against one another in a fixed configuration, while embedded boundaries require Cartesian grids. Neither of these allows relative motion of the cell blocks. Most importantly, none of the methods introduced so far achieves the simplification and efficiency gains that arise from following moving regions’ physics in their own reference frames.

The advantages of working in the most suitable reference frame can be substantial. Consider, for example, a hydrodynamics problem in which structure A, containing only subsonic motions relative to its own center-of-mass that vary on short lengthscales, moves supersonically within a larger background fluid B with longer gradient scales. If such a problem were treated with a moving grid scheme, the time-step within region A would be severely limited by its supersonic velocity and its small cell sizes; transformation to the moving frame could reduce the number of time-steps required by a large factor. Numerical accuracy would also be substantially improved as there would be no need to perform numerous close subtractions of velocities in order to find the relative velocities between cells.

Analogous advantages can stem from equation simplification. Suppose, for example, that in the moving region there is a diffusive transport process that is unimportant in the background. Treating this transport process in the moving frame eliminates what would otherwise be a large, unnecessary, advective flux. If the velocity contrast between the regions approaches the relativistic level, treating everything in the background frame introduces serious conceptual problems because

true diffusion problems cannot be formulated covariantly; transformation to a frame in which local motions are slow permits a clean use of the local Newtonian limit in which diffusive transport is mathematically consistent.

1.3. Our innovation

The system we present here, which we call *PATCHWORK*, is designed to eliminate the limitations due to use of a single reference frame for all patches, while also maximizing its ability for simultaneously dealing with issues of multiple scales, multiple grid symmetries, and multiple varieties of local physics. Exploiting the flexibility of the MPMD approach, it utilizes well-defined coordinate transformations informed by relativistic methods (but not restricted to relativistic problems) to simulate heterogeneous systems in which regions requiring independent treatment are regarded as independent processes operating in independent reference frames. Each patch has its own grid, with its own resolution scale and symmetry; among the many benefits offered by independent local coordinate systems, patches give an easy solution to the problems raised by coordinate singularities (e.g., at the origin of a polar system). Each patch also solves its own equations, whatever is necessary to do the job in that region.

The relationships between these frames are defined by coordinate transformations with the ability to eliminate local mean velocities, so as to reap the benefits just described, while retaining a common overall time so that the entire simulation can advance together. This approach creates an important simplification—the same coordinate transformation that relates array index-space to spatial coordinates can also be used to eliminate bulk velocities. This single coordinate transformation also applies to the metric with respect to the coordinates. In addition, provided that the physical quantities entering into the boundary conditions are scalars, vectors or rank-2 tensors, their transformation from the coordinate system of one region to that of another is well-defined and straightforward. In this way, the free use of arbitrarily curvilinear and arbitrarily discretized coordinates can be combined with the virtues of treating local physics in its most natural reference frame.

However, in order for the patches to exchange boundary conditions at simultaneous times, the time coordinate in all the patches must be the same. In relativistic terms, this means that the coordinate transformations relating patches with relative motion are *not* Lorentz transformations; as a result, the reference frames of moving patches are in general non-inertial. This policy may be somewhat unfamiliar, but because the equations of physics can all be written in completely covariant fashion, their form under this sort of transformation is well-

defined.

Our version of the multipatch system also offers several additional features. Because the separate patches are simulated using independent programs, they can have independent time-steps; because long time-step patches need many fewer updates to traverse the same physical time, when parallelized those patches can be assigned many more cells per processor to achieve better load-balancing. In addition, patches can be added or removed from time to time as conditions change and different demands arise.

To mitigate the complexity and overhead created by inter-patch communications when the computation is parallelized, we have created a client-router-server system that efficiently links the correct processors in each patch to their boundary condition partners in other patches.

Lastly, our package is structured as a “wrapper” to fit around a user-supplied hydrodynamic or magnetohydrodynamic simulation code. The only requirement placed upon these codes is that their primary dependent variables (the ones whose boundary conditions are exchanged between patches) are all the same. In many simulations, they may be the fluid primitive variables, proper rest-mass density, proper internal energy density, and 4-velocity, but that specific list is not a requirement. Thus, different codes can be used in different patches without harm to the system, provided only that they share a common core set of dependent variables.

2. METHOD

The multipatch method is applicable to any numerical simulation code in which the computational domain is discretized into a fixed set of discrete grid-cells, and the primary dependent variables, the ones exchanged between patches as boundary conditions, are the same in every patch. For development and testing purposes, we have used it with the finite volume general relativistic hydrodynamics code *HARM3D* (Noble et al. 2009) running in every patch, but it should be consistent with a great many simulation codes. Subject to the stipulation about variable consistency, different codes can run in different patches.

2.1. Overview

PATCHWORK’s structure is based on the concept of “patches”. A patch is a region of space defined by the user. Locations within it are described by its particular coordinate system and discretized according to its own particular grid. Its physical evolution in time is governed by a particular set of equations, always including the Euler fluid equations, but potentially extensible to the Navier-Stokes equations or the MHD equations, and potentially supplementable by chemical or nuclear reac-

tion networks, a Poisson solver for self-gravity, or other sorts of equations. Evolving an individual patch is the responsibility of an individual process within the MPMD environment. Although it is possible to run the simulation as a single program, using one program for each patch keeps the code simple and conceptually clear. As a result, the method is intrinsically parallelized: there must be at least one processor for each patch.

PATCHWORK coordinates a number of different individual patch processes through the incorporation of several specific routines into the fluid simulation code chosen by the user. One of these routines, the one controlling inter-patch communications, is independent of the individual problem, and is therefore not intended to be altered by users. Others must conform to a standard form, but must be supplied by the user because they are specific to the problem. Not all of these are required for all problems, but three are always necessary. One defines the initial conditions, but differs from conventional initial condition specification only in that the user may need to choose them in a fashion informed by knowledge of how the other patches will be initialized. A second implements the boundary conditions on the outer physical boundary of the problem volume, again an operation occurring in all fluid simulations. The third contains information that is specific to multipatch operation: the coordinate transformation Jacobian between an individual patch’s coordinates and the “background coordinates” (see below for definition); the trajectory followed by the patch relative to the background coordinates; and the location of the problem volume boundaries in background coordinates. Individual patches need to know where the overall problem boundaries are in case they encounter these boundaries as they move. Several other multipatch-specific routines are optional and will be described later.

A system of “background coordinates” underlies the entire region being simulated. The locations and boundaries of all the patches are defined in terms of this system. It is always Cartesian, and its time coordinate is the universal time for all the patches’ coordinate systems. Its purpose is both to serve as a reference for positions and to serve as a “common language” for all patches to describe the locations of exchanged data.

Individual patches can have any shape or size. They can be stationary relative to the background coordinates or move. Their internal coordinate systems and grids are entirely independent of all other patches’ spatial coordinate systems and grids. It is convenient in many problems to divide the patches into two categories, “global” and “local”. Frequently, one patch provides the great majority of boundary condition data for the other patches and occupies all or a large part of the problem volume. When that is the case, that patch is deemed

“global”, and its reference frame is tied to the frame of the background coordinates. Its internal spatial coordinate system, however, can still be defined independently of the background coordinates. All the other patches are then considered to be “local”. Physics in whatever part of the problem volume falling under a local patch is controlled by that patch; when there is a global patch, its process is relieved of all responsibility for computing the time-dependence of the matter within such a covered region.

To make physical sense, there is one smoothly-varying spacetime for the entire problem volume. Its metric, however, can be expressed in terms of any of the specific coordinate systems. Transformations from one coordinate system to another adjust the metric elements accordingly.

Physical consistency likewise demands that all patches are updated according to the same time coordinate and must reach a given value of this time coordinate together. Such synchronization is achieved automatically if all advance with the same time-step. However, as we discuss below (Sec. 2.6), this is not necessary. If some patches can be evolved stably and accurately with a longer time-step than others, it is necessary only for the patches all to be synchronized after one time-step of the patch with the longest step. Note that because there is a single time coordinate for all patches, the coordinate transformations between them are *not* Lorentz transformations unless the relative velocity between the two patches being linked is zero and both patches are inertial.

In the course of each update, patches bordering on one another must exchange boundary condition data. Accomplishing this step is the core of our system.

2.2. Boundary condition exchange between patches

Within any particular patch, we distinguish three types of boundary zones for individual processors. In the first category, there are boundaries whose ghost zones are covered by other processors in the patch. The second category is boundaries lying on the physical boundary of the problem. The third category is of greatest interest to the multipatch scheme, those whose ghost zones are covered by processors assigned to other patches.

The first two can be handled by conventional means. Here we describe how the boundary information is obtained for ghost zones in the third category. We begin by displaying an example so that readers can easily visualize the issues involved (Fig. 1). In this figure, the fluid’s internal energy density is represented by color contours and grid cells are delineated by black lines. We have chosen to follow the fluid mechanics in this example by means of two patches, a finely-resolved Cartesian local patch and a more coarsely-resolved polar global patch

whose radial grid is logarithmically-spaced. In the upper panel of the figure, the physical boundary of the Cartesian patch is shown by the inner white box; the area covered by its “ghost zones”, the cells needed to establish boundary conditions for the physical region, lies between the two white boxes. The lower panel shows the converse situation: the jagged white contour shows the boundary of territory in the global patch not covered by physical cells of the local patch; the cells between that jagged contour and the white cells are where the global patch needs boundary data.

The first step is to discover which processors in which patches have the information. To minimize inter-patch communication time, we organize this process to avoid exchanging unused data. Because this procedure is almost independent of whether the patch needing boundary data is a global or a local patch, for this part of the discussion we call them “patch A” and “patch B”. We begin by labeling all the zones in patch A (here this happens to be the global patch) with an integer array, illustrated in Figure 2. The values in this “flag array” denote whether a zone is a ghost zone, and if so, what type of ghost zone. This array must be updated at each time-step if any of the relevant patches move (at each synchronization time-step in the case of heterogeneous time-steps: Sec. 2.6). In the figure, the white zones are in the interior of patch A, and have nothing to do with boundary conditions. Gray zones are the zones in patch A completely covered by patch B; they, too, are irrelevant to boundary conditions². A zone in the global patch is considered to be covered by the local patch if its center falls within the local patch’s physical region. The red zones are ghost-cells in the patch A grid directly touching physical cells in patch A. Fluxes across their inner (in a topological sense) faces are used in updates of physical patch A cells. Blue zones are the outer layer of ghost-cells; they are used in the internal reconstruction by which cell-center values of physical quantities are extrapolated to the face touching the physical boundary.

It is important to note that this system is thoroughly agnostic about many of the possible choices made in different codes. Because the coordinates at which the boundary data are needed are determined by the fluid code operating in the requesting patch, it doesn’t matter whether that code defines the variables at cell-centers or the centers of cell-faces or anywhere else; it knows the locations at which it needs the information, and it is the job of the responding code (which may be an en-

² If patch A were a local patch, it would have gray zones only if patch B were another local patch, and patch B took priority over patch A; we have not yet implemented “local-local” boundary data exchange, but plan to do so soon.

tirely different one) simply to interpolate its data, no matter how defined in terms of location within cells, to the proper point. The system is even capable of accommodating codes with different numbers of ghost-cells. *HARM3D*, for example, requires three layers of ghost-cells, but *PATCHWORK* contains a parameter that can be set to whatever number of layers the user’s code needs.




Figure 1. A snapshot of internal energy density (color contours) and grid-cells in a 3D blast wave simulation. (Upper panel): White squares show the physical boundary (inner) and numerical boundary (outer) of the local patch. Where the local and global patches overlap, only the local grid is shown. (Lower panel): Like the upper panel, but where the patches overlap, only the global grid is shown. The colored cells inside the jagged loop are filled with data interpolated from the local patch to the global patch. The white cells inside the jagged loop are unused when the patches are in this configuration because the local patch updates the physics in their volume. The physical (inner) and numerical (outer) boundaries of the local patch are shown as thin white lines for reference.

Once patch A determines the locations of its ghost zones’ centers, a list of the background coordinates for




Figure 2. The “flags” assigned to the global grid-cells for the same snapshot shown in Figure 1. Red and blue indicate inner and outer ghost-zones (for the global patch), respectively. White cells are ordinary cells in the global patch interior; gray labels global cells covered by the local patch and ignored. The thin gray squares show the physical (inner) and numerical (outer) boundaries of the local patch.

these locations is sent to all other potential patch B’s along with a request for interpolation values at the coordinate locations. Patch B interpolates within its grid in order to find the values at the locations desired by patch A. It then transforms the data from its coordinate system to the background coordinate system, using the coordinate transformation Jacobian linking patch B to the background system. Only then are the boundary data transmitted back to patch A, which transforms it from the background system to its own coordinates. This procedure enables every patch to deal with the incoming coordinate list independently, without knowing anything about other patches. Doing things this way is especially important when patch B moves.

Note that if patch A is the global patch, the first step is done differently. At initialization, the local patches are informed of the cell locations at which the global patch dependent variable data are defined. Because the local patches know their own positions in terms of background coordinates, they can determine on their own what data the global patch needs. This alternate procedure has the virtue of diminishing inter-patch data transmission.

2.3. Interpolation

Although remarked on only briefly in the preceding sub-section, there are a number of subtleties to data interpolation, and multiple mechanisms may be used. In the current version of our system, we use a comparatively simple method, but this could readily be upgraded to something more sophisticated for problems requiring

it.

In principle, an arbitrary number of zones could be used to support interpolation to a single point. However, it is generally best for the interpolation stencil to extend away from the point by a number of zones that is no more than the number of ghost-zone layers (usually 2 to 3), in order to avoid complications with parallelization of the code.

For our current method, we employ tri-linear interpolation. We locate the grid corner closest to the interpolation point and define the stencil in 3D to be the centers of all eight cells touching that corner. This method works quite well when the dimensions of the cells in patch A and patch B are comparable (see Sec. 3.3), but can lead to errors when they are not. In some sense, this is unsurprising: if there is structure on the finest scale supported by one of the patches, it cannot be well-represented by a much coarser grid in the other. However, the trouble can also move in the opposite direction because the eight cells in the finer grid nearest the interpolation point may together cover only a small part of the volume of the ghost-zone in question if its grid is much coarser. Sometimes errors of this latter variety can be substantially reduced by replacing the values in the inner layer of ghost-cells with a wider average over nearby cells. Such an operation effectively magnifies the volume of the finer-scale grid contributing to the coarser-grid ghost-cells.

Without special methods, interpolation does not necessarily conserve quantities. To achieve strict mass (or momentum or energy) conservation in our data interpolation would require identifying all cells that fall within the ghost-cell and summing their contributions. If the ghost-cell boundaries cut obliquely (or even worse, in a curve) across some of the interpolation cells, one would need to adjust their volumes accordingly. Although this is possible if both global and local patches are in Cartesian coordinates, it becomes a non-trivial mathematical problem once any of the patches are in curvilinear coordinates.

When interpolation that is more nearly conservative is desirable, yet the cell geometry makes rigorous conservation arduous, another avenue may be taken. In this method, one constructs the conserved quantity density for a coarse ghost-cell by averaging the conserved quantity densities for all those cells in the finer-scale patch whose centers are located inside that ghost-cell. This method is not perfect because not all the averaged cells contribute the same volume to the ghost-cell, but when the cell-size contrast is large, there are many small cells entirely contained within the large ghost-cell and only a few at its edges, so the error should be small.

2.4. Adding and removing patches

Stationary or moving patches can be added or removed throughout the simulation anywhere within the physical problem volume. This is done using the flag array for the ghost-zones discussed in Section 2.2. Although these flags are most often used to signal the need for data interpolation from overlapping patches, they can also be used to signal the need to interpolate data for other reasons as well—such as removing or introducing a new patch. To remove a local patch, one temporarily changes the flags on all the zones in the global patch covered by the local patch to “ghost zones” so that all of them are filled with interpolated data provided by the local patch. Once these zones are filled with data, one changes the flags back to their normal state. To add a local patch, one creates a new patch process, and in its initial condition sets all its cell flags to indicate they are ghost zones. Just as in the patch removal operation, these cells are then filled with the data they need, and the flags can be reset to normal as soon as that is done. However, the simulation must be stopped immediately after a patch removal or immediately prior to a patch addition because either one demands a new domain decomposition for parallel processor assignment.

2.5. Parallelization and inter-patch communication

One of the most difficult tasks in developing a multi-patch code is its parallelization. It requires a sophisticated infrastructure combining two levels of data communication. In one, boundary data exchange within a patch, a single executable exchanges information between its multiple processors exactly in the way made familiar by ordinary parallelized methods. In the other level, boundary data exchange between patches, it is necessary to enable effective data communication when the pairing of processors with overlapping boundaries evolves dynamically, and two independent executables, both running within the MPMD environment, must be coordinated.

To describe how we achieve this, we first define a notation. We label each CPU in a simulation by C^i_j , where i is a patch-ID and j denotes local CPU rank within that patch. We set the patch-ID of the global patch to $i = 0$ and that of local patches to $i = 1, 2, \dots, N - 1$, where N is the total number of patches in the simulation, including the global patch. The index j runs from 0 to $n_i - 1$, where n_i is the number of CPUs used for patch i .

Consider a CPU at the edge of a patch, designated C^{i_0} . This CPU possesses boundary zones that need to be filled with interpolated values. It needs to know which CPUs C^{k_j} in other patches are lying under these zones (there could be multiple values of j satisfying this criterion, and sometimes multiple values of k). It must also contact them to request interpolation values. The partner CPUs C^{k_j} , on the other hand, need to

know in advance that other CPUs may be contacting them. Because these relationships constantly change if the patches move relative to one another, this information must somehow be updated dynamically, even though the patches may have differing time-steps.

To solve this problem, we construct a client-router-server system, setting up inter-patch communication relationships that can persist throughout the simulation. In its simplest form, one CPU in each patch is chosen to serve as the router, its liaison with all the other patches. Then, when client CPU C^i_j needs information from beyond the boundary of patch i , it transforms the coordinates of the cell-centers in question to background coordinates and broadcasts that list to processors $C^k_{r_k}$, where the r_k processor in patch k is the designated router for that patch. The router processor in the k th patch then transforms the list from background coordinates to patch k coordinates. If all the cell-centers on the list lie outside patch k , the router replies accordingly. On the other hand, if some of them are inside patch k , the router processor determines which of the other processors working on patch k have responsibility for those cells and distributes the request to those processors. These processors, the servers, interpolate their data to the correct positions, transform the results to background coordinates, and return the results to the router. Finally, the router transmits the information back to the client, CPU C^i_j .

This communication scheme is conceptually simple and easy to code. However, if only a single CPU is given router duties for an entire patch, the communication load is shared very unevenly and the great majority of processors sit idle while waiting for the routers to finish their work. To divide the workload more evenly, we regard *all* processors as potential routers for their patch and redefine the client-router relationship uniquely for each individual CPU (see Fig. 3). These relationships are defined at the beginning of the simulation and remain unchanged unless patches are added or removed. For example, one may decide that C^1_0 always contacts C^0_0 for any information regarding patch 0, C^1_1 always contacts C^0_1 , and so on. The function of the router is unchanged; it still determines which, if any, of the processors on its patch holds the information requested and acts as the go-between connecting clients and servers.

Although the varying numbers of processors per patch make an exactly even division of labor impossible, a simple assignment scheme can spread it in a reasonably even-handed manner. If $C^{p_1}_i$ requires information regarding patch p_2 , it contacts

$$\text{router-}p_2(C^{p_1}_i) = C^{p_2}_{i \bmod n_{p_2}}.$$

Note that CPUs on patch p_2 could have 0 or multiple clients on patch p_1 , depending on their index, n_{p_1} , and

n_{p_2} .




Figure 3. Schematic view of client-router-server relations for the multiple-router scheme. White, blue, and green patches represent the global patch (patch 0), local patch 1, and local patch 2, respectively. Example clients, routers, and servers are marked with C's, R's, and S's, respectively. Data requests are shown by red arrows, data returned by blue arrows. The local patches reside inside the global patch, but they are placed outside the global patch and enlarged for visualization of the information exchange system. The squares in the patches represent CPU domains, not grid-cells. In a possible instance of data exchange, a client CPU in a local patch, C^1_0 , (upper left in patch 1) sends its list of ghost zones to its designated router in the global patch, C^0_0 , (upper left in patch 0). C^0_0 then communicates with appropriate server CPUs on its patch to collect the requested data and returns the data to its client. Simultaneously, CPU C^1_3 also requests data from the global patch, working with its global patch router C^0_3 , which, in turn collects the information from the relevant server CPUs and transmits it back to the client. Even while these two patch 1 CPUs communicate with their partners in the global patch, it is possible for a CPU in the global patch, for example C^0_{15} (lower right in patch 0), to be a client, requesting data from other patches such as patch 2; in this case, the router is C^2_6 .

2.6. Heterogeneous time-steps

One of the common problems in simulating multi-scale systems with grid-based hydrodynamics codes is that the time-step of the entire computational domain is limited to a small value by a few regions with small grid-cells and high characteristic fluid velocity. As a result, the remainder of the simulation, where the intrinsic timescales can be much larger, is required to integrate with unnecessarily short time-steps, leading to a large computational cost. However, the multipatch method, in which different regions are updated by independent processes, allows each patch to have different time-steps while nonetheless evolving the system in a fashion synchronized across all patches. We call this mode of oper-

ation one of “heterogeneous time-steps”, in contrast to the simpler “homogeneous time-step” case in which all patches are forced to have the same time-step.

Heterogeneous time-steps can be managed with great flexibility. The only restriction placed on the time-steps in different patches is that the update times should all be synchronized at intervals equal to the longest of the time-steps, $\Delta T \equiv \max_k(\Delta t_k)$, where, as before, k is an index labeling the different patches. To optimize computational resource use, the user adjusts the number of CPUs assigned to each patch so that the wall-clock time to advance by a time ΔT is approximately the same for all patches.

In practice, the coordination works as follows. At the n th synchronization time t_n^{sync} , the different patches exchange boundary condition data. They then also exchange information about their time-steps so they can determine which is the longest and in which patch it is found (call that patch K). If all the other patches receive their boundary data either from patch K or the problem boundary, the next synchronization time is set to be $t_{n+1}^{\text{sync}} = t_n^{\text{sync}} + \Delta T$. If ΔT is a factor Q_l (> 1 by definition) larger than the time-step Δt_l in some other patch l , patch l performs $\sim [Q_l]$ updates, where $[x]$ is the greatest integer $\leq x$, while patch K works to advance to t_{n+1}^{sync} . When patch l reaches a time t' such that $t_{n+1}^{\text{sync}} - t' < \Delta t_l$, Δt_l is reset to $t_{n+1}^{\text{sync}} - t'$ to achieve synchronization. Patches that have arrived at t_{n+1}^{sync} before the rest of the patches wait until all have reached it. When that has been achieved, the cycle is repeated. Because, by definition, conditions in patch K change by at most a modest amount over a time ΔT , it is unnecessary for the other patches to receive boundary condition information from it during their individual time-steps within the interval ΔT . However, when there are more than two patches, it is possible that some patches may require boundary data from other patches whose time-steps are shorter than ΔT . When that occurs, that pair must exchange boundary data at times determined by the longer of their two time-steps. Note that processors *within* the same patch trade boundary data in the usual way at each of that patch’s internal time-steps.

If all the patches are solving the same equations, optimal load balancing can be achieved when patch K can be identified with reasonable reliability in advance, and the ratios Q_l can similarly be estimated. If those criteria are met, all that is necessary is to assign CPUs in patch K a number of cells $N_K \simeq Q_l N_l$ for each of the other patches labeled by l .

In some situations, the Q_l might be essentially fixed throughout the simulation. For example, this would be the case in a simulation of gas dynamics in an isotropic gravitational potential in which the patches are nested spherical shells. In such a situation, the time-step for

each shell would always be $\simeq [(N_{\phi,k})\Omega(r_{\text{min},k})/2\pi]^{-1}$, where $N_{\phi,k}$ is the number of azimuthal cells in patch k , $\Omega(r)$ is the orbital frequency as a function of radius, and $r_{\text{min},k}$ is the smallest radius in patch k . In such a case, load-balancing could be achieved fairly reliably and would need no adjustment during the simulation.

More often, however, the Q_l may vary as functions of time. When this condition obtains, because the MPMD environment does not permit dynamic domain decomposition, perfect load-balancing will nearly always be an unreachable goal. Nonetheless, as we show in Sec. 4, even approximate load-balancing by combining appropriately chosen numbers of processors per cell in each patch with heterogeneous time-steps can lead to significant gains in computational efficiency. These gains can be sustained even if the ratios Q_l change significantly through the simulation if the user periodically stops the simulation and restarts with an adjusted choice in numbers of processors per patch.

3. PHYSICS TESTS

In this section, we showcase the performance quality of the multipatch method. First we demonstrate that it accurately reproduces the analytic solutions to two classic hydrodynamic simulation test-cases even when critical features of these solutions pass through patch boundaries and the grid symmetries and resolutions of the patches differ sharply. We then explore how well non-conservative mass density interpolation maintains mass conservation and identify the conditions in which it does not.

3.1. Sod shock tube

For this test, we created a square planar 2D problem volume in which, following the Sod prescription (Sod 1978), the fluid is initially at rest everywhere, but there is a sharp pressure and density discontinuity at a specific value of x within the volume. There is no initial variation as a function of y . Within this volume, we placed a square local patch and gave it a constant velocity so that it moves diagonally in the xy -direction. Both patches are described in Cartesian coordinates and have uniform spatial grids in Minkowski space. The local grid was oriented parallel to the global patch grid.

Because HARM3D is framed in terms of relativistic dynamics, it is convenient to choose c as the unit of speed. Given arbitrary code-units of length ℓ_0 and mass density ρ_0 , the unit of time is ℓ_0 and the unit of pressure is $\rho_0 c^2$. To ensure Newtonian flow (Hawley et al. 1984), it suffices to make $p/\rho \ll 1$ when measured in code-units. We also chose an adiabatic index $\gamma = 1.4$.

Our problem volume was 40 code-units on a side, and its 800^2 cells had dimensions 0.05×0.05 ; the local patch had a side-length of 8 and was cut into 320^2 cells of

dimension 0.025×0.025 . For the initial state, the gas was divided at $x = 0$ into left (L) and right (R) states, with density and pressure $\rho_L = 1.0 \times 10^5$, $p_L = 1.0$ and $\rho_R = 1.25 \times 10^4$, $p_R = 0.1$. The sound speed was therefore $\simeq 3\text{--}4 \times 10^{-3}$ on both sides, clearly sub-relativistic. Zero-gradient boundary conditions were used for the problem exterior.

The origin of the local patch was placed initially at $(0, 10)$ in global patch Cartesian coordinates, so that it straddled the discontinuity. Its velocity was $\vec{V} = 1.0 \times 10^{-3}(\hat{x} + \hat{y})$, so that it traveled both subsonically and rather slower than the shock front. As a result, although the shock initially forms inside the local patch, the shock eventually passes out of the local patch. The contact discontinuity stays within the local patch for a longer time, but it, too, eventually leaves it.

In Figure 4, we compare 1D cuts through the data at three different times with exact analytic solutions (Laney 1998). The left-hand column shows $t = 500$, when the shock is completely contained in the local patch; the middle column displays $t = 900$, the time at which the shock emerges from the local patch; the right column illustrates $t = 2500$, when the contact discontinuity is about to leave the local patch. As this figure shows, our numerical results in both the global and the local patches follow the exact solution very closely at all times. Thus, in this test, crossing a patch boundary creates no artifacts, even though the local patch moves obliquely across two different kinds of discontinuity.

3.2. Sedov-Taylor blast wave

The purpose of this test was to show the performance of the multipatch when at least one of the patches has a grid whose symmetry is a poor match to the natural symmetry of the problem. To that end, we study a Sedov-Taylor 3D spherical blast wave (Sedov 1959) with a central local patch using Cartesian coordinates and a global patch using spherical coordinates. As a standard of comparison, we also contrast a monopatch simulation with entirely Cartesian coordinates. Although the Cartesian grids of the local patch are poor matches to the spherical symmetry of the physical problem, they do have the virtue of covering the coordinate singularity at the origin.

A blast wave is formed when a large amount of energy E is deposited in a small region. If the ambient gas is motionless and has uniform density ρ_1 , a spherical shock wave travels rapidly outward. Once the mass swept up by the shock exceeds the mass originally located in the small energy-deposition region, the shock front's radial position as a function of time is given by

$$R_s = \left(\xi \frac{E}{\rho_1} \right)^{1/5} t^{2/5} \quad (1)$$

until E/R_s^3 is small enough to be comparable with the ambient pressure. Here ξ is a dimensionless number ~ 1 .

To simulate this, we follow Fryxell et al. (2000) and divide the initial state into two regions. As in the Sod shock tube problem, we choose the unit of velocity to be c , but use arbitrary code-units for length and mass. In terms of these units, region 1 is a small sphere of radius $\delta r = 25$. Its initial pressure $p_1 = (\gamma - 1)E/(4\pi\delta r^3) = 1$ for adiabatic index γ (again $= 1.4$), while its density $\rho_1 = 10^{-3}$. Region 2 is everything outside $r = \delta r$. Here the initial pressure $p_2 = 10^{-10}$ and initial density $\rho_2 = \rho_1$. The dimensionless coefficient of eqn. 1 is a function of γ ; for $\gamma = 1.4$, it is 1.175 (Ostriker & McKee 1988).

For the monopatch, the computational domain is a cube of dimension 2500 having $N_{\text{mono}} = 400^3$ equal-volume cubical zones with side-length $\Delta = 6.25$. For the multipatch simulation, the global patch is a sphere of radius 1000 described in spherical coordinates, but with two kinds of cut-out: a sphere of radius 290 surrounding the origin and a bi-cone of half opening-angle $\pi/10$ surrounding the polar axis. Its angular grid is uniform with 120 cells in polar angle θ and 320 cells in azimuthal angle ϕ , but its radial grid has 80 logarithmically-spaced cells. The local patch is a cube of side-length 600 centered on the origin with 120^3 equal-volume cubical cells of side-length 5, similar to the Cartesian cell-size in the monopatch simulation. Where the local and global patches overlap, their cell dimensions are likewise similar.

The quality with which all these discretized representations reproduce the analytic result is shown in Figure 5, which again shows the situation at three different times. In this case, at the earliest time the shock front is entirely within the local patch, while it is a short distance outside the local patch in the middle time, and far outside the local patch at the last time. At the earliest time, the data for the Cartesian local patch and the Cartesian monopatch are, not surprisingly, nearly identical; the entire global patch remains in the initial state at this time. Interestingly, the shock at this time is at slightly larger radius than predicted by the analytic solution in both the monopatch and the multipatch simulation. At the middle time, the local patch and monopatch still closely agree, but the shock region is located in the global patch. The global patch data conform noticeably more closely to the analytic solution than do the monopatch data, presumably because the spherical geometry of the global patch better represents the essentially spherical character of the blast wave. This cannot be an ordinary resolution effect because the cell-sizes of the two grids in the shock region are quite similar at this time. At still later times, the monopatch and global patch data are slightly closer to one another, with both showing some departures from

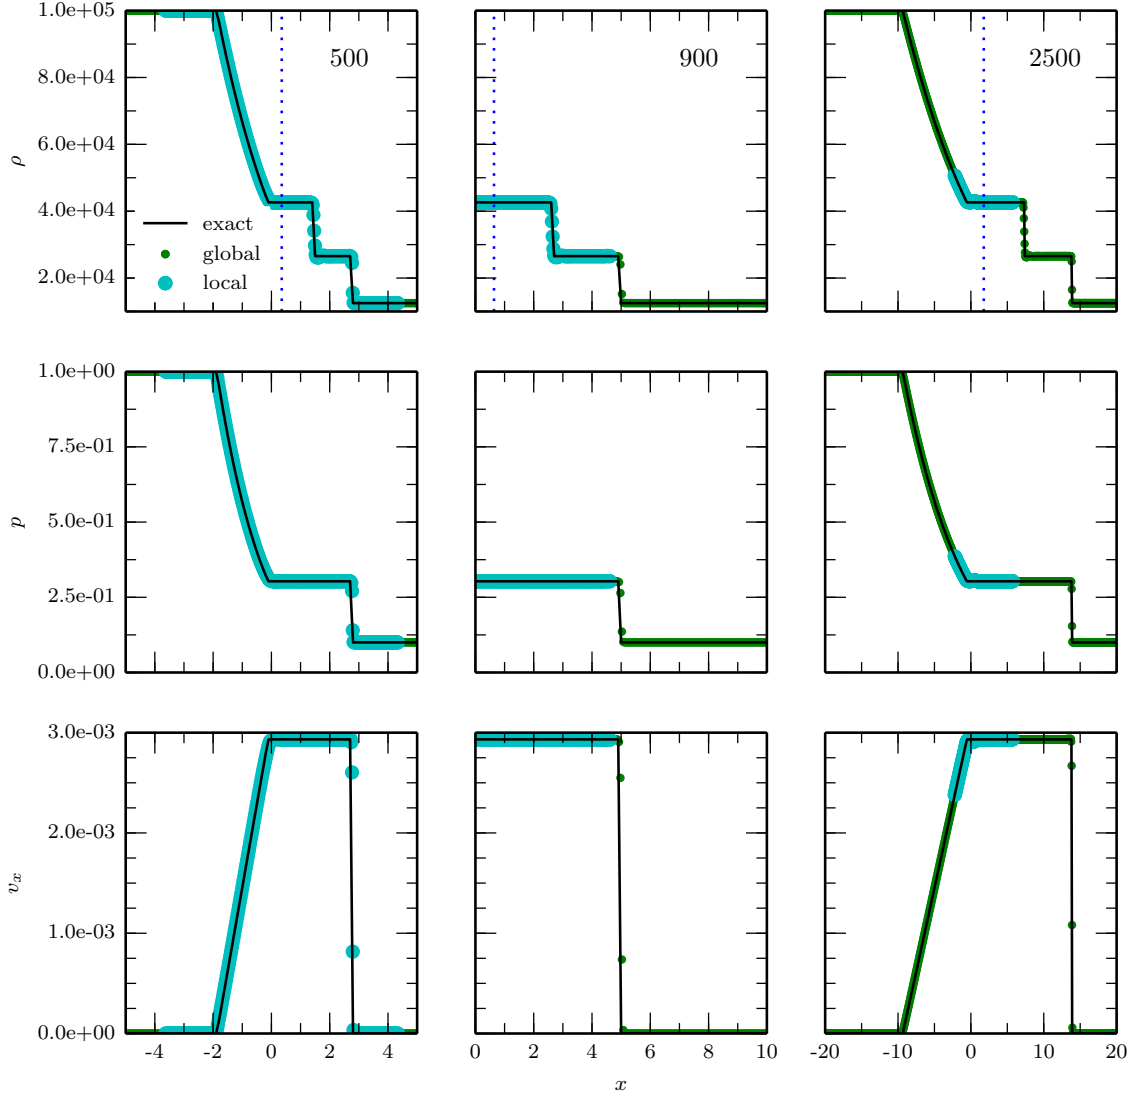


Figure 4. Shock tube test problem in 2D observed at three times, $t = 500$ (left), $t = 900$ (middle), and $t = 2500$ (right). The scale of the horizontal axis is the same for the left and middle columns, but $x = 0$ shifts from the center of the plot (left column) to the left-hand edge (middle column); in the right column, the span of x displayed increases by a factor of 4, but, like the left column, $x = 0$ is in the center of the plot. Each column of three panels shows 1D cuts in density ρ , pressure p , and velocity v_x as functions of x at $y = 10$. Data from the global patch is shown with small green dots, data from the local patch with large cyan dots, data from the analytic solution is shown with a black line. The instantaneous position of the center of the local patch is shown by a dotted vertical line. The shock travels in the $+x$ -direction roughly $3\times$ faster than the local patch moves diagonally toward the upper right, so that it initially forms inside, but then leaves the local patch. Similarly, a contact discontinuity also forms inside the local patch, but likewise eventually leaves it behind.

the analytic solution. We attribute this behavior to the spherical geometry of the global grid compensating for its rather coarser resolution, a result of the logarithmic cell spacing. Most importantly, and shown best by the middle column, it is clear that crossing the patch boundary creates no artifacts because the multipatch simulation gives a very similar result to the monopatch. We therefore conclude that in this sort of situation, departures from the analytic solution are due to the usual limitations of finite discretization, not the multipatch algorithm.

3.3. Inter-patch conservation

As remarked earlier, our interpolation scheme is not strictly conservative, even though many hydrodynamics codes that can be used in concert with our multipatch method are. That contrast makes it worthwhile to examine how large an error may be induced by non-conservative interpolation, and how that error depends on the interaction between problem character and details of multipatch implementation.

In principle, this error could depend on many variables. To simplify the discussion and focus on what we

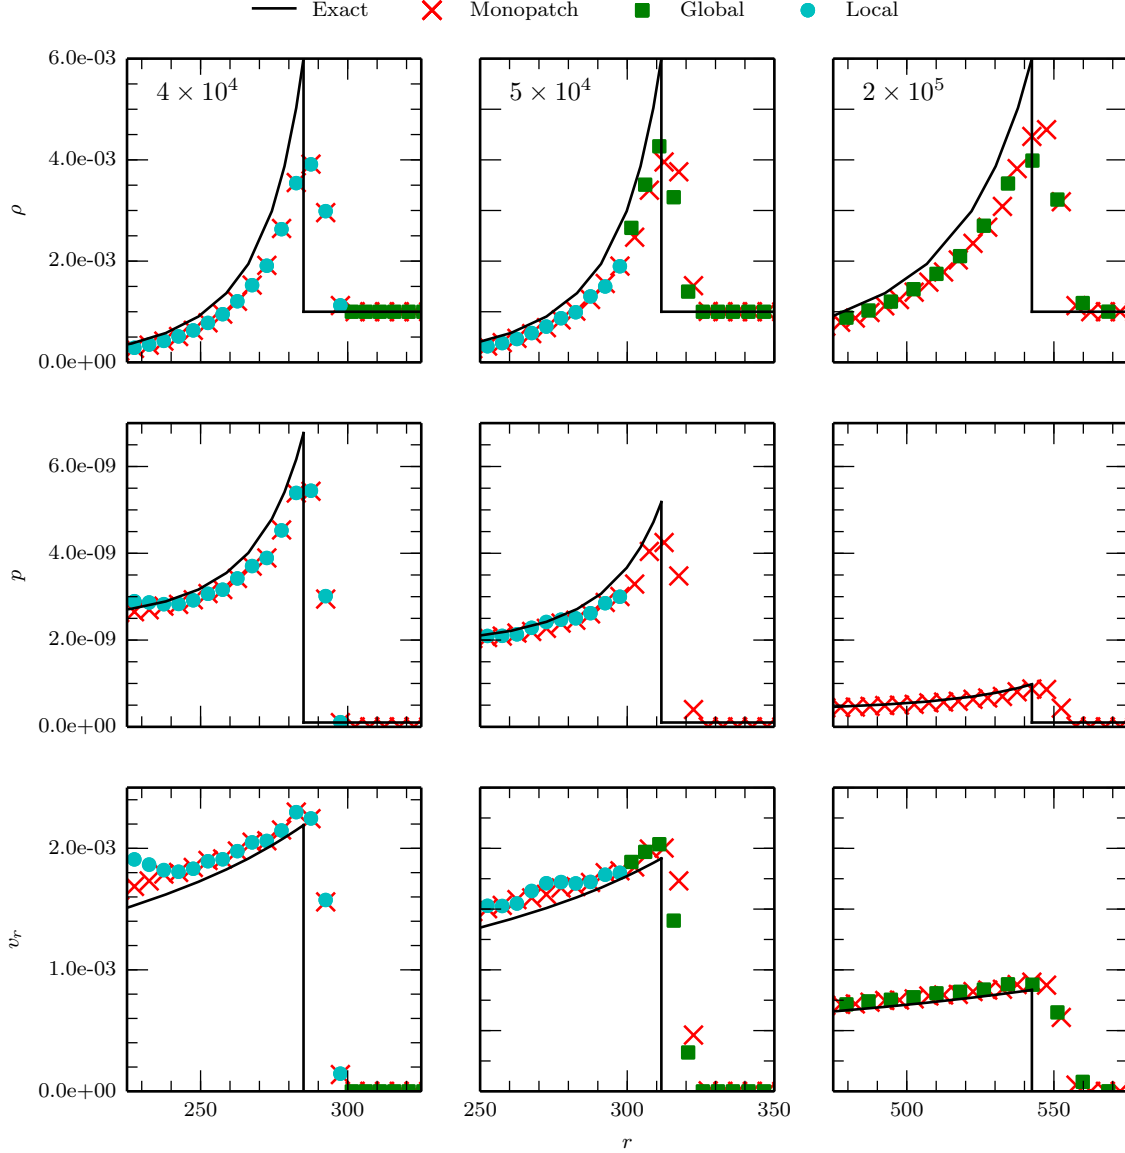


Figure 5. Sedov-Taylor 3D spherical blast wave at three different times ($t = 4 \times 10^4, 5 \times 10^4, 2 \times 10^5$). Monopatch data (red X's) are contrasted with multipatch data (green squares for the global patch, cyan circles for the local patch) and with the analytical solution (black line). Each column of three panels shows 1D radial cuts in density ρ , pressure p , and radial velocity v_r .

believe is the principal issue, we study an idealized problem, one in which matter flows from a patch in which it has acquired order-unity amplitude fine-scale structure into a more coarsely-resolved patch. The parameter that appears to affect conservation errors the most is the ratio between the lengthscale of the structure and the resolution scale of the coarser patch.

To illustrate this dependence, we construct a 3D system in which the problem volume extends from $x = -200$ to $x = +600$ in Cartesian coordinates, but in the transverse directions (y and z) spans only the range $[-20, +20]$. The local patch is stationary and occupies the region $-200 \leq x \leq 0$ in global coordinates.

Both patches have uniform grids that are parallel to each other, but the x -direction cell-sizes of the global patch (5) are $10\times$ that of the local patch (0.5), while the y - and z -direction ratio is 11.4 (5.7 as opposed to 0.5).

The initial condition for the test is shown in Figure 6. At $t = 0$, all of the fluid is traveling at $V_x = 0.1$ in the x -direction, but its density and pressure differ sharply across the line $x = +10$, located a short distance into the global patch from the local patch boundary. To the left of that line, $\rho_L = 1 + \sin^2(\omega_n y)$ and $p_L = 10^{-12}\rho_L$ (as in the previous tests, $c = 1$ in our units), while on the right $\rho_R = 10^{-16}$ and $p_R = 10^{-28}$. This sharp pressure

contrast induces a flow from left to right in the frame of the bulk flow. Because the sound speed on the left is so small ($\sim 10^{-6}$), even in 1000 time-units the high pressure gas expands only a very slight distance to the right in the moving frame of the fluid. Thus, the density and pressure modulation across the patch boundary at $x = 0$ is essentially constant throughout the run of the test.

The frequencies $\omega_n = n\omega_g$ are chosen as multiples of the spatial Nyquist frequency of the global grid resolution in the y direction, $\omega_g = N_{g,y}\pi/L_{g,y}$; this definition also ensures that the total mass in the entire computational domain is the same for each ω_n . The case illustrated in Figure 6 is for the case $n = 0.14$; a close look at the figure shows how the density distribution in the region $0 < x < 10$ is smoothed because the global patch grid is coarser than the local patch grid.

Outflow boundary conditions are enforced at $x = 600$ while reflecting boundary conditions are applied at all other physical boundaries. In the absence of numerical error, these boundary conditions (and the extremely low value of ρ_R in the initial state) guarantee that the total mass on the grid remains constant.

In Figure 7, we show the ratio $\Delta M(t)/M_{\text{flow}}(< t)$ between the change in total mass and the amount of mass that has flowed across the patch boundary up to that time for initial density modulation frequencies $\omega_n = \{0, 0.01, 1.0, 2.0, 2.6\} \times \omega_g$. For all values of ω_n , a small error ($\sim 10^{-6}$ of the mass flow at $t = 1$: a bit larger for $\omega_n = 0.01$ or 0.1 , several times smaller for $\omega_n = 0$) appears immediately. When $\omega_n < \omega_g$, there is no further increase in the absolute mass error for the rest of the test, all the way to $t = 10^3$; as a fraction of the mass flow across the patch boundary, the error therefore decreases $\propto t^{-1}$. Thus, when the receiving grid is capable of representing the modulation, the fractional conservation error is quite small.

However, as perhaps should not be too surprising, when the receiving grid is too coarse to support the modulation, the error is larger. For integer values of ω_n/ω_g , by 10 time-units, the time in which the bulk flow can move to the right by 1 length unit (2 global grid cells), the total mass error becomes proportional to the amount of mass that has passed from one patch to the other, $\Delta M/M_{\text{flow}} \simeq 7 \times 10^{-7}$. For non-integer values of $\omega_n/\omega_g > 1$, the error fraction is smaller, $\simeq 2 \times 10^{-8}$, so the initial error dominates until $t \gtrsim 100$.

In more realistic problems, we have found that dynamical effects can enlarge these errors, and in some cases positive feedback loops can develop. However, the driving factor appears to be similar, the inability of a coarse grid to support variations occurring on lengthscales too small for it to resolve. Devices to curb this sort of interpolation error are problem dependent. We have found

that in some situations, including those in which unstable feedback can occur, it is helpful to apply a smoothing algorithm which averages the hydrodynamic variables of the ghost zones in the global patch after the inter-patch communication. In other circumstances, alternative methods like those described in Sec. 2.3 may be required damp error growth.




Figure 6. Density (grayscale) in the initial condition for the mass-conservation test. Global patch grid lines are given in both patches for reference; note that the y -scale is expanded by a factor of 20 relative to the x -scale, making almost square cells look like narrow rectangles. The stationary local patch covers the global coordinate range $-100 \leq x \leq 0$ for all y . Density values $\lesssim 1$ are white to emphasize the modulation.

4. COMPUTATIONAL EFFICIENCY AND PARALLELIZATION SCALING

Lastly, we present data on computational efficiency and parallelization scaling. When discussing these issues in the context of ordinary monopatch operation, the principal questions generally have to do with the fundamental efficiency of the computational algorithm and the ratio between time spent exchanging boundary condition information and computing updates. The former sets the basic scale in terms of zone-cycles per processor per unit time; the latter is determined by the additional cost incurred by inter-processor communication. When a hydrodynamics code parallelizes well, the fraction of total processor-hours devoted to communication is nearly independent of the total number of processors. Thus, to gauge how much overhead is created by multi-




Figure 7. Change in total mass $\Delta M(t)/M_{\text{flow}}(< t)$ relative to the time-integrated mass flow across the patch boundary for initial conditions defined by frequency $\omega_n/\omega_g = 0$ (black line), 0.01 (blue line and dots), 1.0 (cyan dotted line), 2.0 (yellow dot-dash line), and 2.6 (green dashed line).

patch operation and how efficiently the multipatch system makes use of parallelization, we must contrast multipatch benchmarks with monopatch benchmarks treating the same physics problem, and do so as a function of total number of processors. In addition, we will explore how much our heterogeneous time-step option improves efficiency by contrasting its performance with matched homogeneous time-step runs.

However, all these tests can at best be indicative rather than definitive. Even in monopatch operation, hydrodynamic code speeds can be problem-dependent, and there is every reason to expect that multipatch methods will, if anything, add new ways for code performance to be sensitive to the nature of the specific problem. For example, in multipatch problems the amount of computation required to perform a single zone-cycle can depend on how much effort is necessary to compute coordinate transformations, a quantity that can easily differ substantially from one physical situation to another. If the patches solve different equations, the number of operations per zone-cycle can change even more drastically. Because we expect a significant contrast in overhead between cases in which the local patches either move or are stationary, we will specifically examine that variety of problem-dependence.

To finesse these complexities as best we can, we focus on a single simple test problem: evolving a hydrostatic gas in the absence of any external forces. The background spacetime is therefore Minkowski, and in the initial condition there is uniform density, pressure,

and entropy. The fluid’s adiabatic index is $\gamma = 5/3$.

The problem volume is a 3D cube treated with two patches, a global patch and a local patch. Both use Cartesian coordinates with uniformly-spaced grids. The local patch is a cube with side-length 1/8 the global patch’s side-length, but has a gridscale that is also 1/8 the global patch’s; the two patches therefore have the same number of cells. These choices produce a time-step ratio $\Delta t_g/\Delta t_l = 8$. These will be studied with two different numbers of zones per processor, $n = 20^3$ and $n = 40^3$. Each simulation is run for a fixed time duration, chosen to be just long enough that initialization time is negligible. Zero-gradient boundary conditions are used for the global patch; the local patch never encounters the problem boundary. All cases were performed on the same platform (Texas Advanced Computing Center, Sandy Bridge nodes on Stampede).

In our first set of benchmarks, we consider the case of a stationary local patch and a single time-step for both patches. The results (obtained from STDLIB C “time()”) are shown in the two left-hand panels of Figure 8. In this set, we assign the same number of processors to each patch; that means both patches have the same number of cells per processor. When the number of cells per processor is relatively large, multipatch operations create a very modest overhead: the ratio of cycle-update speed for multipatch to monopatch with 40^3 cells per processor is $\simeq 0.75$. On the other hand, with fewer cells per processor (20^3), the ratio is closer to $\simeq 0.4$.

Another view of stationary patch computational efficiency may be seen in the lower left panel of Figure 8, showing the ratio of time spent in communication relative to the total computational time. Here we define “communication” as any operations involving boundary condition exchange between processors. Examples of communication specific to multipatch operation include determination of client-server relations, transmission of ghost-cell coordinates, and interpolation of data to those coordinates. A fourth category of communication, data transmission from server to client, occurs in any sort of parallelized simulation. “Total time” is defined as all time spent on communication plus time spent on computing hydrodynamic updates; it does not include time spent on ancillary activities such as initialization or writing output. In terms of this measure, we find that for stationary patches the fraction of total time spent in communication is $\simeq 2 - 3\times$ as great as for monopatch runs with the same number of cells per processor. This extra time can be largely attributed to the interpolation step.

In the second set of comparisons (right panels of Fig. 8), we examine what happens when the local patch moves. In this case, the contrast in zone-update rate is

larger and is also a stronger function of the number of processors per patch. As before, larger numbers of cells per processor yield greater efficiency. With 40^3 cells per processor, the update rate for the multipatch is $\simeq 0.7\times$ the monopatch rate for 512 processors per patch, but declines to $\simeq 0.5$ for 1728 processors per patch. The corresponding figures for 20^3 cells per processor are 0.6 and $0.2\times$ the monopatch rate. The fraction of time spent in communication is consistently $5 - 6\times$ larger in moving patches than for the corresponding monopatch case, nearly independent of the number of processors per patch.

The overhead and scaling behavior of moving patches differs from that of stationary patches because, in addition to data interpolation, it is also necessary to determine client-server relationships at each time-step and to transmit fresh ghost-cell coordinate lists. These additional tasks both increase the total overhead and make parallelization scaling poorer. The underlying reason is that as the number of processors per patch grows, more processors must be queried to determine the correct set of client-server connections. The relative load this imposes is larger when the number of cells per processor is smaller because, just like all other boundary data considerations, it is sensitive to the processor domain's surface/volume ratio.

We close this part of the discussion by making an important remark regarding interpretation of these benchmarking results. Although we expect the qualitative trends to be robust, their quantitative character is sensitive both to the specific problem and to the specific architecture of the computing system used. Different problems (and different algorithms applied to the same problem) can have different numbers of arithmetic operations per cell-update, while different cluster architectures can give different rates of inter-processor data transmission. Because part of the multipatch overhead depends on additional data transmission, the relative speeds for these two sorts of processes can alter the multipatch/monopatch comparison at a quantitative level. As we discuss below (Sec. 5), technical improvements in implementation of the multipatch method can also lead to quantitative changes in efficiency.

Next, we compare the computational expense for a multipatch program running with and without the heterogeneous time-step algorithm at fixed total number of zones in each patch N . The time-step ratio between the two patches is $\Delta t_g/\Delta t_l \sim 8$. For ideal load balancing with a heterogeneous time-step, the global patch should therefore have 8 times as many zones per processor as the local patch. In Figure 9, we show how the number of processor-hours required to reach a fixed physical time depends on the number of zones per patch for both a homogeneous and a heterogeneous time-step. As one

might expect, both scale very closely to linearly with the number of zones per patch, but the heterogeneous time-step requires a number of processor-hours $2 - 3\times$ smaller than homogeneous time-step operation.

Note that the gains achieved by use of heterogeneous time-steps appear in a way quite distinct from typical gains in efficiency. They do not lead to any increase in zone-cycles per processor per second; instead they lead to a decrease in the total number of zone-cycles required to accomplish the simulation. In this way, their effect resembles the use of non-uniform spatial grid cells, a device leading to economies in the total number of zones computed by concentrating them where they are most needed, rather than a conventional increase in computing efficiency.

5. CONCLUSIONS

We have presented the essential methods underlying our implementation of a new multipatch infrastructure, *PATCHWORK*, designed to support multiscale, multiphysics, and multireference frame fluid simulations. This method offers a number of advantages for the numerical study of complex fluid problems involving subregions with contrasting properties. Each patch can have its own coordinate system and spatial grid, differing in geometry and resolution from all the other patches (one of the many ways this can be useful is that if the coordinate system preferable for a part of the problem contains coordinate singularities, they can be covered with a new patch). If different regions demand contrasting time-steps, the independence of the processes evolving the patches permits them to have separately-determined time-steps, potentially saving significant amounts of computing. Although the method assumes that a fluid exists throughout the problem volume, if different auxiliary processes are important in different regions (e.g., chemical reaction networks or self-gravity), their patches can treat those processes without burdening the other regions. Lastly, but possibly most importantly, substructures within the problem may have differing preferred reference frames; these, too, can be accommodated easily.

The patches are linked to one another solely through boundary condition exchange. Contrasting grid systems are reconciled through interpolation; contrasting grid geometries and reference-frames are reconciled through coordinate transformations and ensuring that all transformed physical quantities are well-defined scalars, vectors, or tensors.

Parallelization is essential to modern large-scale computing. Arranging exchange of boundary condition information between the correct processors can be a complex problem in a multipatch system when the patches move relative to one another. We have con-

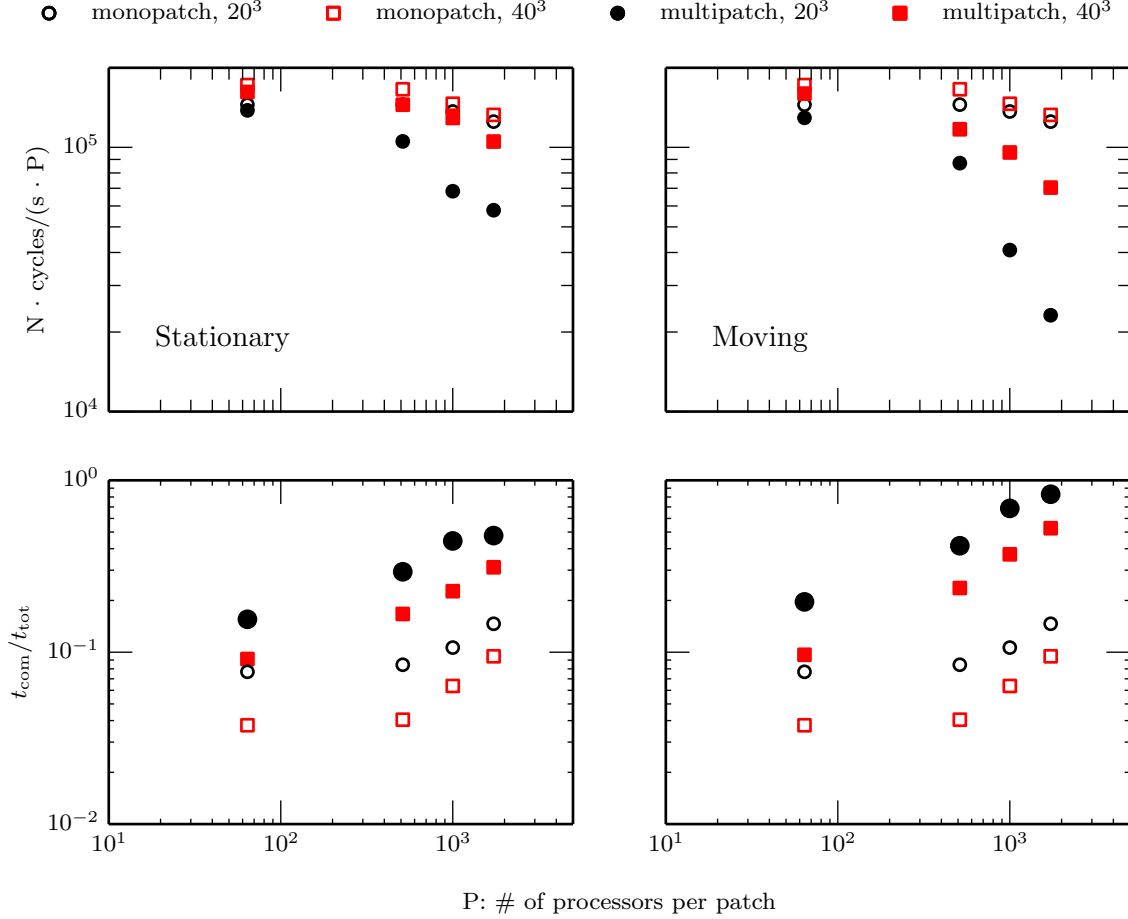


Figure 8. Computational efficiency as a function of numbers of processors per patch and for different numbers of cells per processor. Monopatch method data are plotted with open symbols, multipatch with filled symbols. Runs with 20^3 cells per processor are shown with black circles, runs with 40^3 cells per processor with red squares. Left (right) panels show multipatch simulations with a stationary (moving) patch. Top panels: Processing speed in zone-cycles per processor per second. Bottom panels: Fraction of total wall-clock time spent on communication.

structed a solution to this problem—a client-router-server framework—that updates these connections efficiently. When the patches are stationary relative to one another, the connections need to be identified only once, so the overhead due to multipatch operations is fairly small, especially for larger numbers of cells per processor. When they move, the overhead is more significant and scales with the number of processors per patch, producing a reduction in cell-update rate of about a factor of ~ 1.4 for 512 processors per patch or a factor of ~ 2 for 1728 processors per patch when using 40^3 cells per processor. We note, however, that these comparisons assume that monopatch and multipatch approaches use the same total number of cells; because multipatch operation permits tuning the grid to match local requirements, in practise multipatch simulations may use a much smaller total number of cells than would be required for a monopatch simulation of the same problem—if a monopatch simulation could deal

with the problem at all.

Many extant fluid codes are automatically consistent with this infrastructure. Its sole substantive stipulation is that the dependent variables involved in boundary data exchange should be consistent in all patches. Although we were motivated to build this system by relativistic problems, and our transformation methods are familiar because of their frequent application to relativistic dynamics, in fact they really stem from more general considerations of differential geometry; they therefore apply to any context in which scalars, vectors, and tensors can be defined.

PATCHWORK may be refined and extended, both in terms of its computational efficiency and the span of physical problems on which it can be used. The amount of time spent on interpolation and inter-patch data transmission can be reduced by minimizing the number of arrays transferred or by eliminating unnecessary steps in the coordinate transformations. Mov-




Figure 9. Computational expense of a multipatch simulation with a homogeneous (blue circles) and heterogeneous (orange X's) time-step for a given total number of zones per patch N .

ing from an MPMD environment to one in which a single program employs task-based parallelization will permit dynamical processor reassignment, amplifying the economies in total zone-cycles necessary to compute that accrue from the use of heterogeneous time-steps. Task-based parallelization may also improve interpolation efficiency because, for any single time-step, only a fraction of the processors assigned to an individual patch are involved in inter-patch data exchange. Another improvement will be to add interpolation options that, over a

broader range of circumstances, more nearly conserve quantities that should be conserved such as mass and momentum. Given suitable patch resolution, our current default method does not create significant errors, but it would be valuable to create new schemes, more nearly conservative, that would permit greater freedom in resolution choices. Similarly, some special devices will be necessary to extend our multipatch method to MHD problems in a way that preserves divergence-free magnetic field. We are currently developing algorithms to achieve this and hope to report on them in the not-too-distant future.

ACKNOWLEDGMENTS

We thank Matt Duez for an insightful conversation at the onset of our work. R. M. C., J. H. K. and H. S. were partially supported by NASA grant NNX14AB43G and NSF grant AST-1028111. R. M. C. and J. H. K. also received support from NSF grant AST-11516299. S. C. N. received support from NSF grants AST-1028087, ACI-1515969, and AST-1515982.

This research project used computational resources at the Maryland Advanced Research Computing Center (MARCC). We also thank the NSF for providing XSEDE resources on the Stampede cluster through allocation TG-MCA95C003. Additional resources were provided through Blue Waters sustained-petascale computing NSF projects ACI-0832606, ACI-1238993, OCI-1515969, and OCI-0725070. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications.

REFERENCES

- Adams, M., Colella, P., Graves, D., et al. 2014, Lawrence Berkeley National Laboratory Technical Report, LBNL-6616E
- Arcavi, I., Gal-Yam, A., Sullivan, M., et al. 2014, *ApJ*, 793, 38
- Auchettl, K., Guillochon, J., & Ramirez-Ruiz, E. 2016, *ArXiv e-prints*, arXiv:1611.02291
- Barney, B. 2017, *Introduction to Parallel Computing*
- Berger, M. J., & Colella, P. 1989, *Journal of Computational Physics*, 82, 64
- Berger, M. J., & Oliger, J. 1984, *Journal of Computational Physics*, 53, 484
- Blakely, P. M., Nikiforakis, N., & Henshaw, W. D. 2015, *A&A*, 575, A103
- Brown, D. L., Chesshire, G. S., Henshaw, W. D., & Quinlan, D. J. 1997, in *SIAM conference on parallel processing for scientific computing*, Los Alamos National Lab., NM
- Clough, K., Figueras, P., Finkel, H., et al. 2015, *Classical and Quantum Gravity*, 32, 245011
- Duffell, P. C. 2016, *ApJS*, 226, 2
- Duffell, P. C., & MacFadyen, A. I. 2011, *ApJS*, 197, 15
- Fryxell, B., Olson, K., Ricker, P., et al. 2000, *ApJS*, 131, 273
- Gezari, S., Heckman, T., Cenko, S. B., et al. 2009, *ApJ*, 698, 1367
- Grinberg, L., Fedosov, D. A., & Karniadakis, G. E. 2013, *Journal of Computational Physics*, 244, 131
- Hawley, J. F., Smarr, L. L., & Wilson, J. R. 1984, *ApJ*, 277, 296
- Laney, C. 1998, *Computational Gasdynamics* (New York: Cambridge University Press)
- Mayama, S., Tamura, M., Hanawa, T., et al. 2010, *Science*, 327, 306
- Mocz, P., Pakmor, R., Springel, V., et al. 2016, *MNRAS*, 463, 477
- Nie, X., Robbins, M. O., & Chen, S. 2006, *Physical Review Letters*, 96, 134501
- Noble, S. C., Krolik, J. H., & Hawley, J. F. 2009, *ApJ*, 692, 411
- Ostriker, J. P., & McKee, C. F. 1988, *Reviews of Modern Physics*, 60, 1
- Pollney, D., Reisswig, C., Schnetter, E., & Diener, P. 2017, *llamacode.org*
- Schnetter, E., Blazewicz, M., Brandt, S. R., Koppelman, D. M., & Löffler, F. 2014, *ArXiv e-prints*, arXiv:1410.1764
- Schnittman, J. D. 2013, *Classical and Quantum Gravity*, 30, 244007
- Sedov, L. I. 1959, *Similarity and Dimensional Methods in Mechanics*
- Sod, G. A. 1978, *Journal of Computational Physics*, 27, 1
- Springel, V. 2010, *MNRAS*, 401, 791